Claims

[c1] 1. A system for automatically determining a set of at least one maximal utility object from a set of at least one object represented by a probabilistic model, comprising:
determining an upper bound for the utility of each object;
sorting the objects by the upper bounds in order of highest to lowest;
obtaining a set of known object values for a particular entity;
using the probabilistic model in combination with the information known about an entity to begin predicting the set of maximal utility objects from the set of objects;
examining the utilities associated with each object in the set of objects in the sorted order for selecting maximal utility objects until the set of maximal utility objects us full; and
continuing the examination of utilities until the utility associated with a lowest utility object in the set of maximal utility objects is greater than the upper bound of the utility of a next sorted object in the set of objects.

[c2] 2. The method of claim 1 wherein the entities represent users and wherein the objects represent possible user choices.

[c3] 3. The system of claim 1 further comprising a utility function for adjusting the upper bounds of the utilities.

[c4] 4. The system of claim 3 wherein the upper bounds of the utilities are weighted via the utility function.

[c5] 5. The system of claim 4 further comprising a user interface for weighting the upper bounds of the utilities.

[c6] 6. The system of claim 3 wherein different weights are assigned to upper bounds of the utilities associated with different objects.

[c7] 7. The system of claim 3 wherein different weights are assigned to upper bounds of the utilities associated with different groups of objects.

[c8] 8. The system of claim 3 wherein at least one of the upper bounds of the utilities having a highest value is assigned a penalty for reducing the magnitude of the upper bound.

[c9] 9. The system of claim 3 wherein the objects are resorted by the upper bounds from highest to lowest after adjusting the upper bounds.

[c10] 10. A computer-readable medium having computer executable instructions for dynamically extracting at least one highest probability object recommendation from a probabilistic model without examining all possible probabilistic recommendations from the probabilistic model, said computer executable instructions comprising:
extracting upper bounds of particular states of probability distributions for objects represented by the probabilistic model;
sorting the upper bounds in order of highest to lowest;
examining objects represented by the probabilistic model in order of the sorted upper bounds for each object for determining at least one highest probability object recommendation; and
terminating the examination of the objects as soon as the upper bound of the lowest probability recommended object is greater than an upper bound of a next sorted object.

[c11] 11. The computer-readable medium of claim 10 further comprising a user interface for defining a utility function for adjusting the upper bound of the particular states of the probability distribution for at least one of the objects.

[c12] 12. The computer-readable medium of claim 10 further comprising a user interface for assigning at least one unique utility function for individually adjusting each of the upper bounds of the particular states of the probability distributions of at least two sub-groups of the objects.

[c13] 13. The computer-readable medium of claim 10 further comprising using a set of user preferences in combination with the probabilistic model for determining the at least one highest probability object recommendation.

[c14] 14. The computer-readable medium of claim 10 wherein the probabilistic model is a decision tree.

[c15] 15. The computer-readable medium of claim 10 wherein the probabilistic model is a neural network.

[c16] 16. The computer-readable medium of claim 10 wherein the probabilistic model is automatically generated using machine learning techniques.

[c17] 17. A method for determining at least one highest probability recommendation from a probabilistic model, said model representing at least one object using a probability distribution for representing each object, comprising:
determining an upper bound of a particular state of the probability distribution representing each object;
sorting each object represented by the probabilistic model by sorting the upper bounds associated with each object in order of highest to lowest;
determining a set of user preferences for a particular user;
extracting at least one highest probability recommendation from the probabilistic model based on the set of user preferences for the particular user, wherein the objects represented by the model are examined in the sorted order for extracting the at least one highest probability recommendation; and
terminating the examination of the objects and the extraction of highest probability recommendations as soon as a lowest upper bound of any of the highest probability recommendations is greater than an upper bound of a next sorted object.

[c18] 18. The method of claim 17 wherein the set of user preferences is determined implicitly.

[c19] 19. The method of claim 17 wherein the set of user preferences is determined explicitly.

[c20] 20. The method of claim 17 wherein a utility function is used to adjust the

upper bounds associated with at least one of the objects represented by the probabilistic model.

[c21] 21. The method of claim 17 wherein at least one of the objects having the highest upper bounds are assigned a penalty for reducing a magnitude of those upper bounds so as to reduce a probability of those objects being recommended.

[c22] 22. The method of claim 21 wherein the penalty is user adjustable.

[c23] 23. The method of claim 20 wherein the utility function is adjustable via a user interface.